Blockfrost's Transformation to Not-for-Profit — Voting Rationale

Governance Voting Rationale	
GAID	*gov_action12sumv9qky4pkenmqp7gshv9nxqdk9zyn8gkd77fewkyy3t5tnxtsq457vgq*
Title	Blockfrost's transformation to not-for-profit
Type of GA	Treasury Withdrawal
Amount	₳9,832,979 (~$1,868,266 at $0.19/₳) · 18 months · staffing 79.1% · ops/infra 19.3% · legal/accounting 1.6%
Proposer	Input Output (IOG)
Date submitted	*Epoch 639 (Jun 23, 2026)*
Expiration Date	*Epoch 646 (Jul 28, 2026)*

Contents

- 1.0 Introduction
  - 1.1 Summary
  - 1.2 Description of Governance Action
- 2.0 Discussion
  - 2.1 Method
  - 2.2 The sort, published
  - 2.3 The good, stated as a relation
  - 2.4 The wall, and the entrenchment test
  - 2.5 Reversal, and the transfer that inverts leverage
  - 2.6 The floor, and the six rights
  - 2.7 The private-surplus burden

  - 2.8 What the vote does not reach
  - 2.9 Fiduciary verification (external instrument)
- 3.0 Conclusion, and what would make this a YES

1.0 Introduction

1.1 Summary

We are voting **NO** on this governance action. The NO is **remediable, not terminal** — the register is instructional, and this section is written to be read alongside §3.0, which names precisely what would make a resubmission earn a YES.

The distinction matters here more than usual, because the thing this proposal is trying to do is right. Blockfrost is the access layer through which a majority of Cardano's transactions are submitted and the great majority of its developers read the chain, and it currently sits inside a single commercial entity. Recognising that as a structural risk and moving to place it under community ownership is exactly the kind of response the derivation would ask for. The failure the proposal names is real, and the *kind* of remedy it reaches for — a commons genuinely constituted rather than a commons enclosed — is the correct kind.

The NO is about the structure of the remedy, not its intent. As drafted, the action transfers critical-infrastructure IP — source, trademarks, domains — irreversibly into a newly created entity *before* the governance that is supposed to protect the commons has been elected or proven; it concentrates stewardship of that commons in a board four of whose five seats are reserved for infrastructure companies, several of whom are candidate commercial beneficiaries of the entity's own contemplated future; it lets the entity report on its own governance integrity through its own dashboard; and it reframes a public-good grant as an investment expected to return profits to the treasury. Each of these is a property of *this* action, each fails a published test, and each is fixable. A proposal that fixed them would not be a different idea — it would be this idea, structured so that it guards the failure it exists to prevent rather than reconstituting it in a distributed form.

Because the defects are structural and the transfer they concern is irreversible, the honest token is a NO that deploys against *this structuring* while naming the structuring that would earn assent — not an abstention, which would lower the passage bar for exactly the irreversible move most in need of scrutiny. An external fiduciary read, run independently, reaches the same conclusion from the capital-allocation side.

1.2 Description of Governance Action

This is a Treasury Withdrawal of ₳9,832,979 (~$1.87M at a reference rate of $0.19/₳) funding an 18-month transition of Blockfrost from an IOG-owned service into a free, community-governed public API operated by a not-for-profit. Blockfrost provides a hosted REST abstraction over chain data and transaction submission; the proposal reports it as the number-one hosted platform in the ecosystem (71.5% developer adoption in 2025), serving ~1.84 billion monthly requests, with more than half of all transactions in most epochs submitted through it, and roughly 90% of its traffic on the free tier.

The action would establish a not-for-profit (newly formed, or hosted within an existing organisation such as PRAGMA), governed by a five-seat board — four seats reserved for open-source infrastructure entities, one community seat — elected on-chain by end of Q1 2027, with a named preliminary board seating it in the interim and IOG holding a non-voting advisory seat for the first year. All Blockfrost IP would transfer to the entity. The budget is 79.1% staffing (a team of six, salaries reported as averages), 19.3% ops and infrastructure (~$20k/month), and 1.6% legal and accounting. Administration runs through Intersect with a third-party assurer and the Sundae Labs treasury smart-contract framework, with a public dashboard, quarterly reports, a 99% uptime SLA, proportional refund on reduced scope, and unspent funds returned. Long-term sustainability — including a possible commercial tier with profits routed to the treasury, or a vendor-backed partner model — is explicitly deferred to the elected board.

2.0 Discussion

2.1 Method

Every governance action is sorted for depth before its merits are read and before any standing policy is applied. The sort returns a lane — how much analysis the action earns and which instruments run on it — never a verdict. The token is selected downstream, under the abstention-spine discipline, where the operative question is not whether a vote helps a proposal pass but whether a warrant has been derived to deploy delegated stake against it. The sort is published so that a reader who disagrees can locate whether the disagreement is about the sort or about the analysis; those are different arguments, and only one is about the proposal. For an action that builds a standing institution, the ordering is not a formality — it is what keeps the institution's evident usefulness from pre-deciding the structural questions.

2.2 The sort, published

	Reading
Axis 1 — wall / knob / wall-building	**Wall-building.** The action does not tune a setting; it constitutes a new standing institution — a not-for-profit owning critical infrastructure, a reserved-seat board with an ongoing mandate over it. We considered and declined a Gate 2 arrest: that fail-state is reserved for a structural default installed through an instrument that *does not acknowledge* it, and this proposal openly declares the institution it builds. But the vehicle is a spending ballot, and the object is a permanent steward of the majority access layer — an instrument-fit gap carried below as a Spine 1 flag. Wall-building routes to the full lane and faces the entrenchment test in the confirming direction.
Axis 2 — exit-remediability	**R1, decaying toward R0.** Measured at the failure mode, not the money. The funding is R3 — sunsets, refunds, unspent returns. But the failure mode is a captured or drifting entity that now *owns* the infrastructure. Before transfer the community holds leverage (it funds the transition); after the IP transfer and board constitution, leverage inverts and there is no stated claw-back. Sort on the terminal band.
Axis 3 — epistemic-dependency depth	**D2 operationally; D3 on governance integrity.** A majority of transaction submission and most developer reads condition on this layer. More sharply: the entity reports on its own governance integrity through its own dashboard, and the parties positioned to capture it are also the parties reporting on it — a reading that conditions its own detection. A clear semiotic signature accompanies it: the reframing of a public-good grant as an investment returning profit.
Lane	**Full lane + mandatory distal sensing + mandatory reflexive check + entrenchment test.** The deepest routing the instrument produces.

The sort earns its keep here by refusing the surface. By sector and ask-shape this reads as an infrastructure grant; three tiers of scrutiny separate it from one. What the axes catch is that this is not maintenance of a good but the constitution of the body that will own it, transferred irreversibly, depended on by a majority of the network, and reporting on itself.

2.3 The good, stated as a relation

The relation is real and it is large. Blockfrost is the maintained capacity to read from and write to the chain without operating one's own node — for most of the ecosystem's builders, and for the submission of a majority of its transactions. That capacity anchors settlement access (Right I, at the submission path specifically) and self-determination (Right VI, the ability to build without a rent-extracting gatekeeper). Kept free and non-enclosed, it is a genuine public good.

But the question the intake form asks — the good as a *relation* rather than a deliverable — surfaces that this proposal funds two goods, not one, and they sort differently. The first is the ongoing free public API: non-rival, non-excludable at the free tier, covering mainnet, preview, and preprod. The second is the *governance structure* that will own and steward it. The first is a commons relation. The second is an instrument-layer object — a standing administrator of critical infrastructure — and it is where the analysis concentrates, because the durability of the first depends entirely on the design of the second. A free public API owned by a body that can later enclose it is a free public API on a term the body sets.

2.4 The wall, and the entrenchment test

Wall-building is not presumptively good and not presumptively bad; it faces the entrenchment test in the confirming direction. A proposed wall must show that the failure it guards is genuinely irremediable or self-concealing — otherwise it is a knob welded shut, and a knob welded shut is a standing invitation to capture on the far side.

On the first half of the test the proposal does well: the failure it guards — a single commercial entity controlling the majority access layer, able to enclose it by raising prices or removing the free tier — is genuinely high on dependency and genuinely hard to reverse once enclosed. That is a failure worth building a wall against.

On the second half the proposal does not yet pass. The test is not whether *a* wall is justified but whether *this* wall guards the named failure. As designed, it risks reconstituting the enclosure it exists to prevent, in distributed form. Four of five board seats are reserved for infrastructure companies; the entity's own contemplated sustainability path adds commercial tiers and paid vendor memberships; the parties who would hold those seats are candidate beneficiaries of that path. The wall meant to end single-entity control of the access layer, built this way, could hand stewardship to a consortium of vendors with a live interest in the commercial surface — trading one enclosure risk for another and calling the trade decentralisation. Operationally the proposal *does* decentralise, and impressively — a hundred-plus Icebreakers, a federated operator model. But governance is being centralised into one entity and one board at the same moment operation is being spread out, and it is the governance layer, not the operator layer, that decides whether the good stays free.

2.5 Reversal, and the transfer that inverts leverage

The decay flag is the heart of the matter. At the moment of the vote the community has maximal leverage: it holds the funds the transition needs. The proposal spends that leverage to transfer the IP — source, trademarks, domains — into the new entity, and names no condition under which the transfer reverses. After enactment the leverage is gone and the entity owns the asset. If the board drifts commercial, aligns with vendor interests, or simply governs the commons poorly, the community's recourse is to contest a standing entity that now holds the infrastructure a majority of the network depends on — R1 at best, and decaying, because each operating cycle deepens the dependency and entrenches the entity. "Unspent funds returned" governs the money; nothing governs the asset. The absence of an IP reversion or claw-back tied to defined governance-integrity or free-tier conditions is the single most consequential gap in the proposal, because it is the gap that makes every other gap irreversible.

2.6 The floor, and the six rights

The Hippocratic floor asks whether the action, at minimum, does no harm to the commons relations it touches. Read on intent, it clears easily; read on structure, it does not clear cleanly, and the framework reads structure. The finding is a fitness gap, not an accusation: the action, as drafted, leaves material and inadequately controlled risk to the very relations it means to protect.

Right V — commons integrity — is the central right and the one with the most bite. The contestability test (can the arrangement be challenged and replaced?), the enclosure-trajectory test (does this trend toward enclosure?), and the reversibility test (can it be undone?) all read poorly against an irreversible IP transfer to a vendor-weighted entity with a contemplated commercial layer. Right III — governance participation — is touched by the board's concentration: a 4:1 reservation of stewardship of a commons asset to infrastructure companies is a concentration finding on its face. Right I — settlement access — is protected now by the free tier but placed at future risk by the contemplated commercial tiers and the absence of any binding free-tier covenant surviving to the elected board. Right IV — informational integrity — is met for uptime and usage, which register externally on-chain and on the dashboard, but not for governance integrity, which the entity reports on itself; and the investment-return framing engages the narrative-accuracy test. Right VI — self-determination — tracks Right I: protective if the free good holds, at risk if vendor membership fees reconstitute the gatekeeper. Only Right II is largely untouched, save that the "investment returning profit to treasury" language brushes the store-of-value reframing the ontology flags.

Several rights, then, are touched, and several in a corrosive-leaning or at-risk direction. This is not the clean protective profile of a maintenance grant. It is a genuinely mixed reading in which the good is real and the structure imperils it.

2.7 The private-surplus burden

The burden governs any treasury withdrawal and asks a question of claim structure, not of the recipient: is a commons relation maintained, stated in relation terms, or is a private surplus transferred? The recipient's non-profit form does not answer it — a non-profit funding an activity whose benefit is captured privately does not discharge the burden, and the framework does not ask the recipient's tax status.

The proposal *attempts* the burden and articulates the relation intelligibly: the free public API, the majority submission path, the developer onboarding it enables. That is real and it is to the proposal's credit. But the discharge is contested at more than the margin, for two reasons. First, the sustainability path introduces a private-surplus vector the proposal does not foreclose — commercial tiers and vendor membership fees flowing to board-adjacent companies, with the profits-to-treasury return stated only as something the board "could consider," which is to say not binding. Second, the proposal reaches heavily for exactly the arguments the burden says do not discharge it: the KPI section argues TVL, MAU, transaction volume, and protocol revenue, and the sustainability section reframes the grant as an investment expected to return profit. These answer a quantity question — *will the number rise* — in place of the relation question, and the burden is explicit that once that substitution is accepted the legitimacy question has been answered in the quantity's terms before it could be asked. Growth is answerable specifically: growth reflecting real productive activity serves the relation; growth in value locked, in revenue, or in attention does not. The proposal shows productive activity in its usage data, but rests its case on the disallowed frame.

2.8 What the vote does not reach

The full lane runs distal sensing and the mandatory reflexive check; several findings belong in the trajectory log rather than in the vote, and one is a reflexive disclosure owed to delegators.

The reflexive check (mandatory at D3). This DRep holds only the general participant interest — like everyone, our governance practice and chain reads may pass through the access layer this action concerns. General interest is the condition of participation and does not disqualify; we hold no specific interest, sit on no board, and benefit from no vendor structure. Carried plainly, as the discipline requires.

The proposal's own reflexivity — logged, diagnostic. The preliminary board seats vendors (BlinkLabs, TxPipe, Sundae Labs, Masumi/Begin) who are candidate beneficiaries of the vendor-backed model the same document proposes, and Sundae Labs appears in three roles at once — preliminary board member of the entity, member of the Oversight Committee verifying this withdrawal's administration, and author of the treasury smart-contract framework holding the funds. None of this is alleged as bad faith; it is a concentration and contestability signal (Right III, Right V) that a redesign should manage rather than merely disclose.

The semiotic and temporal signature. The investment-return reframing, and the sequence it opens — fund the transition, seat a vendor-weighted board, let the board add commercial tiers, let vendors become paid partners — is a trajectory in which no single step is the enclosure but the sequence is. This is precisely the aggregate-temporal-semiotic harm signature distal sensing exists to read, and it is why the reversibility gap in §2.5 is not a technicality.

The instrument mismatch (Spine 1, carried not cast). The ballot approves a spend; the object is the constitution of a standing steward of critical infrastructure. A revision might reasonably split these — a governance action ratifying the entity and its mandate with binding covenants, and a funding action sized to the transition — so that the community's approval of money is not silently read as approval of an institution and its seat structure. We carry this as a finding rather than a token because the action is not merely a trajectory signal; it has fixable defects the ballot *can* reach, and reaching for abstention where a remediable NO is owed is the rationalisation the spine's own guard warns against.

Portfolio note. IOG and affiliates disclose ₳130.7M in treasury allocations across projects to date. Disclosed and to the proposal's credit; logged as a concentration datum for the portfolio view the vote itself does not take.

2.9 Fiduciary verification (external instrument)

This DRep's framework reads the relational and structural question and by design does not read the capital-allocation question — price, instrument fit, size-calibrated discipline, upfront exposure, opportunity cost. Under our own field-fitness discipline an unexamined region raises vigilance rather than lowering it, so we ran the proposal through the [DRep Treasury Rule Book v17] (credit:

@InputEndorsers) as a fiduciary gate — one that can lower a merits-YES but cannot raise a merits-NO.

Scorecard selection is itself a judgment here and we state it as one: the primary value is transitioning critical infrastructure into a community public good, which routes to the Public-Good and Civic-Service scorecard with the governance-work and recurring-maintenance modules active — but the contemplated commercial arm is a material future workstream that, under the mixed-proposal rule, would have to pass its own gates and is presently unscorable because it is deferred to the future board. That deferral is itself a finding.

The gate does not clear, and it does not clear independently of the relational read. Classified **Very Large** by nominal request, the proposal faces the highest non-systemic bar — an exceptional public return, low upfront exposure, stress testing, and a clear opportunity-cost case — and several of those are absent: there is no risk register, no stress testing of the capture and commercial-drift scenarios, and no binding return. More decisively, the rulebook's own override discipline maps a wrong or contested *instrument* and a failed *enforceability* gate to No or defer regardless of score, and both fire here: the self-contained-withdrawal preflight is weak (entity form, board, IP mechanics, and sustainability model are all to-be-determined-later), asset-recovery and continuity rights are absent (the enforceability gate), the conflicts around the vendor board are disclosed but unmanaged (the governance module's independence red flag), and the sustainability story leans on the self-financing promise the maintenance module explicitly red-flags.

The substantive contribution is a convergence at the level of the conclusion, not merely the soft spot. From the capital-stewardship side, the gate reaches the same four defects the relational read reached from the rights side — the un-clawed-back transfer (enforceability), the vendor-board concentration (neutrality and conflicts), the deferred self-financing sustainability (the maintenance red flag), and the investment reframing (instrument fit and public return). Two instruments built on unrelated foundations arriving independently at No/defer is the strongest form the gate's corroboration takes, and it sharpens rather than merely seconds the resubmission specification below.

3.0 Conclusion, and what would make this a YES

We are voting **NO — remediable.** The idea is right and the ecosystem plainly wants it; the structure, as drafted, transfers a critical commons irreversibly into a vendor-weighted, self-reporting entity before the governance meant to protect it is proven, and rests its treasury case on an investment frame the burden does not admit. The register is instructional, and the token deploys against this structuring rather than against the transition itself, because the transfer at its centre is the kind that cannot be undone by inaction once made.

A resubmission would earn a YES by structuring the wall to guard the failure it names. Concretely, the specification is: bind the free public good with a covenant that survives to the elected board and cannot be removed without a defined community process; attach an IP reversion or claw-back triggered by defined governance-integrity or enclosure conditions, so the transfer is not the point of no return; rebalance the board so stewardship of a commons asset is not four-fifths reserved to companies that are candidate commercial beneficiaries, and manage — not merely disclose — the vendor and multi-role conflicts; provide governance-integrity verification that does not run through the entity's own reporting; replace the investment-and-KPI case with the relation case the burden asks for, showing productive activity rather than projected return; and consider splitting the institution's ratification from its funding so the ballot that approves money is not read as approving a permanent structure. Each item is reachable, and together they would convert this from a wall that risks reconstituting enclosure into one that ends it.

The vote answers the action; the log answers the trajectory. This action, as structured, we cannot assent to — and the structuring that would earn assent is within reach, which is why the NO names the path rather than closing it.

DRep ID: drep1yfaq8dsam7nusdccey2x2p684f6ulhr42pv24tslv0terqs3nq50q

DRep Profile: [DRep Talk Profile]

[The Coordination Commons] - DRep reference

Stay in touch. X: https://x.com/styg50

[IPFS of rationale]